

This listing of claims will replace all prior versions, and listings, of claims in the application.

Listing of Claims:

1. (Currently Amended) A server system comprising:

a high performance spinning-type hard drive for storing a first set of data; and

a power controlled spinning-type hard drive for storing a second set of data,

wherein an entirety of the power controlled spinning-type hard drive is configured for lower power consumption as compared to the high performance spinning-type hard drive and the first set of data is distinguished from the second set of data by characteristics, the characteristics are determined by the server which accesses content policies to evaluate the data before determining whether the data is stored in the high performance hard drive or in the power controlled hard drive.

2. (Original) The server system of claim 1, wherein the power controlled spinning-type hard drive achieves a lower power consumption by spinning up the power controlled drive upon receipt of a service request and spinning down the power controlled drive after providing service.

3. (Original) The server system of claim 1, wherein the first set of data comprises data that is requested at a rate above a predetermined threshold, and the second set of data comprises data that is requested at a rate below the predetermined threshold.

4. (Original) The server system of claim 1, wherein the first set of data comprises data that has a last accessed date after a predetermined date, and the second set of data comprises data that has a last accessed date before a predetermined date.

5. (Original) The server system of claim 1, wherein the first set of data comprises data that has one of a creation date and a modification date after a predetermined date, and the

second set of data comprises data that has one of a creation date and a modification date before the predetermined date.

6. (Original) The server system of claim 1, wherein data is placed on the power controlled spinning-type hard drive after classification of the data into one or more of the characteristics of a fully distinguished file name, file size, creator, owner, consumer, groups, distribution lists, access control list detail, certificates, signature attributes, protocols, content resolution, encoding technique, encryption technique, key properties, internal subjects, keywords, content tags, assemblies, associations to other files, replication, caching, directory and related database extensible properties.

7. (Original) The server system of claim 1, wherein the first set of data and the second set of data comprise audio or video files.

8. (Original) The server system of claim 1, wherein the first set of data is comprised of replicas of World Wide Web documents not cached on high performance storage by other World Wide Web servers and the second set of data is comprised of World Wide Web documents which are cached on high performance storage by other World Wide Web servers.

9. (Original) The server system of claim 1, wherein the high performance drive spins continuously.

10. (Original) The server system of claim 1, wherein a subset of the first set of data is moved to power controlled drive in accordance with a predetermined condition.

11. (Original) The server system of claim 10, wherein the predetermined condition is based on at least one of the last access date of the subset of data, the current date, the number of times the subset of data has been requested in a period of time, the data size, users of the data, a file name, the data types, and internal content.

12. (Original) The server system of claim 1, wherein a subset of the second set of data is moved to the high performance drive in accordance with a predetermined condition.

13. (Original) The server system of claim 12, wherein the predetermined condition is based on at least one of the last access date of the subset of data, the current date, and the number of times the subset of data has been requested in a period of time.

14. (Currently Amended) A method for managing data on a server, comprising:

classifying the data to be managed into a first set and a second set;

storing a first set of data on a high performance spinning hard drive associated with the server; and

storing a second set of data on a power controlled hard drive, wherein an entirety of the power controlled hard drive is configured for lower power consumption as compared to the high performance spinning hard drive and the first set of data is distinguished from the second set of data by at least one of characteristics of the data and characteristics of metadata, and wherein the characteristics are determined by the server which accesses content policies to the data before determining whether the data is stored in the high performance hard drive or in the power controlled hard drive.

15. (Original) The method of claim 14, wherein the classifying comprises analyzing the data according to one or more of the characteristics of a fully distinguished file name, file size, creator, owner, consumer, groups, distribution lists, access control list detail, certificates, signature attributes, protocols, content resolution, encoding technique, encryption technique, key properties, internal subjects, keywords, content tags, assemblies, associations to other files, replication, caching, directory and related database extensible properties.

16. (Original) The method of claim 14, further comprising:

determining the first set of data and the second set of data by comparing the request rate of data to a predetermined threshold;

assigning data that has a request rate above the predetermined threshold to the first set of data; and

assigning data that has a request rate below the predetermined threshold to the second set of data.

17. (Original) The method of claim 14, further comprising:

determining the first set of data and the second set of data by comparing the last request date of data to a predetermined date;

assigning data that has a last request date after the predetermined date to the first set of data; and

assigning data that has a last request date before the predetermined date to the second set of data.

18. (Original) The method of claim 14, further comprising:

determining the first set of data and the second set of data by comparing the creation date of data to a predetermined date;

assigning data that has a creation date after the predetermined date to the first set of data; and

assigning data that has a creation date before the predetermined date to the second set of data.

19. (Original) The method of claim 14, wherein the first and second sets of data comprise audio or video.

20. (Original) The method of claim 14, wherein the first and second sets of data comprise documents.

21. (Original) The method of claim 14, further comprising spinning the high performance drive continuously, and maintaining the power controlled drive in a reduced performance state unless actively accessing data stored thereon.

22. (Original) The method of claim 14, further comprising moving a subset of the first set of data to the power controlled drive in accordance with a predetermined condition.

23. (Original) The method of claim 22, wherein the predetermined condition is based on at least one of the last access date of the subset of data, the current date, and the number of times the subset of data has been requested in a period of time.

24. (Original) The method of claim 14, further comprising moving a subset of the second set of data to the high performance drive in accordance with a predetermined condition.

25. (Original) The method of claim 24, wherein the predetermined condition is based on at least one of the last access date of the subset of data, the current date, and the number of times the subset of data has been requested in a period of time.

26. (Currently Amended) A computer-readable medium having computer-executable instructions for performing a method of managing data on a server, the method comprising:

classifying the data to be managed into a first set and a second set;

storing a first set of data on a high performance spinning hard drive associated with the server; and

storing a second set of data on a power controlled hard drive, wherein an entirety of the power controlled hard drive is configured for lower power consumption as compared to the high performance spinning hard drive; and wherein the first set of data is distinguished

from the second set of data by at least one of characteristics of the data and characteristics of metadata, and wherein the characteristics are determined by the server which accesses content policies to evaluate the data before determining whether the data is stored in the high performance hard drive or in the power controlled hard drive.

27. (Original) The computer-readable medium of claim 26, wherein the step of classifying comprises analyzing the data according to one or more of the characteristics of a fully distinguished file name, file size, creator, owner, consumer, groups, distribution lists, access control list detail, certificates, signature attributes, protocols, content resolution, encoding technique, encryption technique, key properties, internal subjects, keywords, content tags, assemblies, associations to other files, replication, caching, directory and related database extensible properties.

28. (Original) The computer-readable medium of claim 26, further comprising the step of spinning the high performance drive continuously, and maintaining the power controlled drive in a reduced performance state unless actively accessing data stored thereon.

29. (Original) The computer-readable medium of claim 26, further comprising the step of moving a subset of the first set of data to the power controlled drive in accordance with a predetermined condition.

30. (Original) The computer-readable medium of claim 29, wherein the predetermined condition is based on at least one of the last access date of the subset of data, the current date, and the number of times the subset of data has been requested in a period of time.

31. (Original) The computer-readable medium of claim 26, further comprising the step of moving a subset of the second set of data to the high performance drive in accordance with a predetermined condition.

32. (Original) The computer-readable medium of claim 31, wherein the predetermined condition is based on at least one of the last access date of the subset of data, the current date, and the number of times the subset of data has been requested in a period of time.
33. (Original) The system of claim 1, wherein a database is divided such that a first portion of the database, having a first set of attributes, is stored in the high performance hard drive, and a second portion of the database, having a second set of attributes, is stored in the power controlled hard drive.
34. (Previously Presented) The system of claim 33, wherein the first set of attributes comprises a first set of meta-directives and the second set of attributes comprises a second set of meta-directives.
35. (Original) The system of claim 1, wherein the high performance hard drive comprises one of a remote storage area network (SAN) device and a network attached storage (NAS) device, wherein power management directives control the storage of data across multiple storage devices.
36. (Original) The system of claim 1, wherein the power controlled hard drive comprises one of a remote storage area network (SAN) device and a network attached storage (NAS) device, wherein power management directives control the storage of data across multiple storage devices.
37. (Original) The system of claim 1, wherein the power controlled hard drive comprises one of a remote storage area network (SAN) device and a network attached storage (NAS) device, wherein power management directives are used to control power management events in storage devices.
38. (Original) The system of claim 1, wherein the power controlled hard drive is a redundant array of independent disks (RAID).

DOCKET NO.: MSFT-2933/306027.01
Application No.: 10/823,918
Office Action Dated: November 17, 2006

**PATENT
REPLY FILED UNDER EXPEDITED
PROCEDURE PURSUANT TO
37 CFR § 1.116**

39. (Original) The system of claim 38, wherein the second set of data is distributed across the redundant array using a stripping algorithm.
40. (Original) The system of claim 39, wherein the stripping algorithm reduces the power up impact to a smaller set of drives upon data access operations.